arXiv:1612.02575v1 [cs.CV] 8 Dec 2016

Filter sharing: Efficient learning of parameters for volumetric convolutions

Rahul Venkataramani*, Sheshadri Thiruvenkadam, Prasad Sudhakar
Hariharan Ravishankar, Vivek Vaidya

GE Global Research, Bangalore, India

December 9, 2016

Abstract

Typical convolutional neural networks (CNNs) have several millions of parameters and require a large amount of annotated data to train them. In medical applications where training data is hard to come by, these sophisticated machine learning models are difficult to train. In this paper, we propose a method to reduce the inherent complexity of CNNs during training by exploiting the significant redundancy that is noticed in the learnt CNN filters. Our method relies on finding a small set of filters and mixing coefficients to derive every filter in each convolutional layer at the time of training itself, thereby reducing the number of parameters to be trained. We consider the problem of 3D lung nodule segmentation in CT images and demonstrate the effectiveness of our method in achieving good results with only few training examples.

1 Introduction

Deep learning using CNNs are quite successful in learning tasks due to its ability in exploiting spatial context and weight sharing between pixels. With tens of convolutional layers in a typical network, the number of parameters to be learned runs into hundreds of thousands and therefore requires huge amount of training data. In cases of problems involving medical images, annotated data is hard to come by and the dimension of clinical data nowadays, being high-resolution 3D/4D, is huge. Deep learning on such data would require 3D/4D convolutional filters to properly capture the spatial context. In view of above, we run into problems of overfitting and poor generalization, reducing confidence for clinical deployment. Secondly, utilizing large CNN models in devices is also limited by the memory requirements. Often, good CNN models take up several gigabytes

*The first three authors contributed equally. Corresponding author: sheshadri.thiruvenkadam@ge.com

of memory and it is hard to deploy on small devices. In this context, it is of great interest to reduce the number of parameters of a CNN model to make it trainable using small data and also amenable for easier deployment.

It is observed that CNN filters show significant redundancy in the construction of the output maps [8]. E.g., output maps often use scaled versions of some filters. In this work, we exploit this redundancy to reduce the effective number of parameters without altering the network architecture and performance. We learn only a small set of filters for each layer and then learn how to combine these filters to arrive at the full set of filters.

The issue of reducing model size and complexity has been receiving a lot of attention of late. Most of the methods advocate learning full set of filters first and then pruning them. In [10], a method for compressing the model parameters is investigated, along with a scheme to fine tune the network to recover some loss incurred due to compression. A similar approach is presented in [2] where the low-frequency nature of the filters is exploited. An elaborate scheme to quantize and Huffman code the parameters is presented in [7]. To efficiently perform the arithmetic operations during the test time, [4] uses low-precision fixed point representation of the model parameters.

A few methods also exploit the low-rank nature of the convolution filters to efficiently perform convolutions and also to reduce the storage requirements [6, 9]. A few methods address the problem of directly reducing the number of training parameters. [8] exploits the separability property of filters to reduce the number of trainable parameters and [5] presents a method to predict the parameters during training by exploiting their the redundancy.

A work that is in the same spirit as ours is reported in [11]. They propose to use the scheme in [5] during training and then perform a decomposition of these learned filters using a small set of filters, whereas we propose to learn the small set of filters and the mixing coefficients during training time itself by performing back-propagation directly on the filter bank and the mixing coefficients.

2 Learning a Linear decomposition of CNN filters

Consider a convolutional layer in a deep CNN which takes in N D-dimensional feature maps $\boldsymbol{F}_j, 1 \leq j \leq N$ and outputs M D-dimensional feature maps

$$\boldsymbol{G}_i = \left(\sum_{j=1}^{N} \boldsymbol{F}_j * \boldsymbol{v}_{ij} \right) + \boldsymbol{b}_i, \ 1 \leq i \leq M, \tag{1}$$

where $\boldsymbol{v}_{ij}$ is a filter of size $S^{\frac{1}{D}}$ in each of the D-dimensions, $*$ is D-dimensional convolution and $\boldsymbol{b}_i$ is bias. The CNN learns the filters $\boldsymbol{v}_{ij}$ and the bias $\boldsymbol{b}_i$ by using training data, a loss function and back-propagation algorithm. The number of unknowns to be learned in each layer, sans the bias term, is MNS, which can get very large when we have $D > 2$. Filter sharing overcomes the burden of large number of parameters by deriving the MN filters using a smaller set of

P ($\ll MN$) filters

$$\boldsymbol{v}_{ij} = \sum_{p=1}^{P} \widetilde{\boldsymbol{v}}_{\boldsymbol{p}} \alpha_{ij}^{p},\ 1 \leq i \leq M, 1 \leq j \leq N, \tag{2}$$

where $[\alpha_{ij}^{1}, \alpha_{ij}^{2}, \cdots, \alpha_{ij}^{P}]^T$ is a vector of weights which combines the P filters to produce $\boldsymbol{v}_{ij}$.

Parameter reduction can also be achieved by reducing the number of filters in each layer. However, this will fail to capture the necessary filter diversity leading to poor performance. Unlike methods which first learn MN filters during training and then trim them down to a smaller set of filters for the test time, we propose to directly learn the filters $\widetilde{\boldsymbol{v}}_p$ and the weights α_{ij}^{p} explicitly during the training time itself. We call this approach as *filter sharing* method and with this approach, the number of parameters to be trained becomes $MNP + PS$. With filter sharing, the number of parameters scales by P when the filter size increases, whereas otherwise it would scale by $M \times N$, which can be very large. We experimentally demonstrate the filter sharing application in the next section.

3 Experimental results

In this section, we demonstrate the efficacy of our approach on a challenging 3D medical imaging segmentation problem. We first use CIFAR-10 as a test-bed dataset to demonstrate the usefulness of our method on limited data problems.

Re-writing the filters $\boldsymbol{v}_{ij}$ as a linear combination of the seed filters $\widetilde{\boldsymbol{v}}_p$, as in Eq. (2) can be seen as an ill-posed matrix factorization problem. Thus, as a means of regularization of the seed filters $\widetilde{\boldsymbol{v}}_p$ and mixing coefficients α_{ij}^{p}, a variety of schemes were experimented including unit norm constraint on the seed filters, sparsity, rank minimization norms on of α_{ij}^{p} and dropout on the elements of the output feature maps $\boldsymbol{G}_i$. Dropout regularization with lesser dropout probablity than for regular CNN implementations gave the best results compared to the above possible regularization schemes.

3.1 CIFAR-10

Using the CIFAR-10 dataset, we wanted to investigate the robustness of proposed approach to limited training data and also to show its applicability on a non-medical, 2-D labeling task. Towards this goal, we define a customized/exemplar architecture shown in Fig. 1a. We would like to highlight that the goal of this experiment is not to advance state-of-the-art results on CIFAR-10 dataset but to demonstrate that proposed workflow can be incorporated to any CNN architecture without adversely affecting performance. Fig. 1b shows the results of the exemplar architecture on validation dataset of 10000 CIFAR-10 samples learnt on varying sized subsets from $50k$ training samples. Filter sharing achieves an average $\sim 3-4\%$ increase in validation accuracy for all the subsets. Additionally, CIF-CNN-FS achieves 80% validation accuracy using only $\sim 35\%$ of the training dataset ($\sim 17k$ samples) while without filter sharing $\sim 33k$ samples are required to achieve the similar performance.

Layer	# of filters
Conv1	32
Conv2	32
Conv3	64
Conv4	64
Fully connected layer	512

(a)

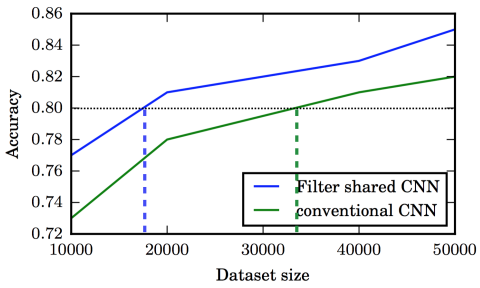

(b)

Figure 1: (a) Exemplar architecture (CIF-CNN) and (b) CIFAR-10 results on validation set. CIF-CNN was used with $P = 15$ in each layer for the filter sharing results in (b).

3.2 Lung nodule Segmentation

Segmentation of relevant structures from 3-D volumetric medical data is critical to morphology measurements for anatomies, characterization of lesions and other applications. We have chosen lung nodule segmentation from 3-D Low Dose CT (LDCT) images as an application for the proposed filter sharing method due to its clinical relevance and challenging nature. We work with 3-D LDCT volumes from Lung Image Database Consortium [1]. The LIDC-IDRI database consists of data from 1,010 patients along with annotation by four experienced radiologists. We pre-selected 93 large solid data with ~ 280 lesions, of which 50% of lesions (~ 140) were used for training, and remaining 50% was equally split for validation and testing purpose.

Fully convolutional networks (FCN) have recently emerged as state of the art algorithms for natural and medical image segmentation, multi-object scene labeling and related tasks. Additionally, FCNs have analysis-synthesis arms to produce output volumes same/similar to the size of input volumes leading

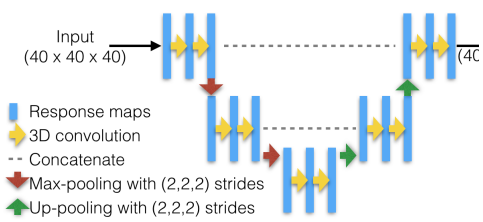

(a)

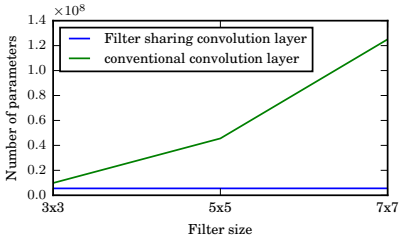

(b)

Figure 2: (a) 3D U-Net architecture and (b) Number of parameters to be trained as a function of the size of convolution filter for 3D U-Net architecture

Architecture	Number of parameters	Dice overlap in %
3D U-Net without FS	9847689	60
3D U-Net with FS	**5505572**	**66**

Table 1: Effect of using filter shared 3D convolutional layers in U-Net

to roughly double the number of parameters compared to a standard CNN label prediction network, making them an ideal candidate to demonstrate the proposed method's usage. Fig. 2a shows 3-D U-Net architecture we have used in our experiments [3]. Fig. 2b shows how the number of parameters in a 3D U-Net convolution layer scales with respect to the size of convolutional filter. As the size of the filters increases, filter sharing method clearly shows its advantage. We would like to highlight that our task is segmentation of nodules and not detection which we assume is done automatically or manually. Input to the architecture are $40 \times 40 \times 40$ LDCT volumes pre-identified as containing nodules and output are the segmentation masks predicted by the network.

We use Dice overlap between the predicted nodule segmentation and the ground truth annotation (provided by an expert) as the metric to evaluate performance. Table 1 contains the comparisons. We achieve a significant 6% increase in average Dice overlap, in addition to achieving 45% reduction in parameters using FS. We hypothesize that learning lesser parameters in limited data scenarios, like medical imaging problems, can lead to more efficient, stable and better architectures.

4 Discussion

In this paper, we have proposed a method to reduce the number of parameters to be learned in a CNN by linearly expanding each convolutional filter in terms of a small set of filters. This method has great advantages when training data is small to avoid overfitting. Moreover, reduction in the number of parameters also helps in reducing the overall memory footprint of the model, making it amenable for deployment in diverse situations.

With filter sharing framework, we can also impose regularization on the kind of filters to be learnt. Also, we can include a fixed set of standard filters (e.g., wavelet filters) to improve the performance. Learning a latent set of filters also helps in easily transferring models across different deep learning architectures. This is extremely relevant when we would like to arbitrarily increase or decrease the number of filters in a given layer of a deep network. These issues are a part of our future work.

References

[1] S. G. Armato III, G. McLennan, L. Bidaut, M. F. McNitt-Gray, C. R. Meyer, A. P. Reeves, B. Zhao, D. R. Aberle, C. I. Henschke, E. A. Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.

[2] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen. Compressing convolutional neural networks. *CoRR*, abs/1506.04449, 2015.

[3] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432, 2016.

[4] M. Courbariaux, Y. Bengio, and J-P David. Training deep neural networks with low precision multiplications, 2015. ICLR workshop.

[5] M. Denil, B. Shakibi, L. Dinh, M. Ranzato, and N. D. Freitas. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156. 2013.

[6] E. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. *Exploiting linear structure within convolutional networks for efficient evaluation*, volume 2, pages 1269–1277. NIPS, 2014.

[7] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *ICLR*, 2016.

[8] Y. Ioannou, D. P. Robertson, J. Shotton, R. Cipolla, and A. Criminisi. Training CNNs with Low-Rank Filters for Efficient Image Classification: ICLR 2016 Poster, May 2016.

[9] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. In *British Machine Vision Conference*, 2014.

[10] Y-D. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. *CoRR*, abs/1511.06530, 2015.

[11] A. Sironi, B. Tekin, R. Rigamonti, V. Lepetit, and P. Fua. Learning Separable Filters. *IEEE Trans. on PAMI*, 37(1):94–106, July 2014.